



Published in final edited form as:

Theor Criminol. 2018 May 1; 22(2): 258–278. doi:10.1177/1362480617707948.

Can we avoid reductionism in risk reduction?

Seth J Prins and

Columbia University, USA

Adam Reich

Columbia University, USA

Abstract

Risk assessment and risk reduction have become increasingly central to criminal justice policy and practice in the last 25 years. Yet there remains a lack of consensus both on the theoretical and methodological foundations of risk and on its social and practical implications. Some proponents see risk assessment and reduction as solutions to the inefficiencies and injustices of contemporary mass incarceration. Some critics see actuarial risk as being partially *responsible* for mass incarceration, and warn that recent iterations will only reinscribe existing inequalities under a new guise of objectivity. Both perspectives contain elements of truth, but each falls short because neither adequately specifies the different dimensions of risk that condition its effects. Using two prominent frameworks as foils, this article excavates the contested terrain of risk assessment and exposes a set of distinctions that can inform the use—and prevent the abuse—of risk knowledge in criminal justice policy.

Keywords

Criminal justice; criminology; methodology; risk assessment; risk reduction

For at least the past 25 years, criminal justice theorists have observed the growing importance of risk assessment and risk reduction in criminal justice policy and practice. Yet there remain conflicting perspectives on how “risk” should be understood and implemented in the criminal justice field (Garland, 2003a; Hannah-Moffat, 2012). Among its proponents, risk assessment seems to offer a solution to the inefficiencies and injustices of mass incarceration (Sherman, 2007). Focused-deterrence and hot-spot policing strategies, for example, promise to concentrate resources on the riskiest people and places while lessening the negative impacts of aggressive policing overall (Braga and Weisburd, 2010; Kleiman, 2009). Allocating probation supervision and treatment resources according to criminogenic risk profiles focuses attention on those most likely to benefit from it, while leaving alone those whom supervision might harm (Andrews and Bonta, 2010; Lowenkamp et al., 2006).

Among its critics, however, risk profiling—through the idea of “selective incapacitation” popularized in the 1980s—is central to understanding the *rise* of mass incarceration and its

Reprints and permissions: sagepub.co.uk/journalsPermissions.nav

Corresponding author: Seth J Prins, Department of Sociomedical Sciences, Columbia University, 722 W 168th Street, 548, New York, NY, 10032, USA. sjp2154@columbia.edu.

Author Manuscript

detrimental effects (Feeley and Simon, 1992; Harcourt, 2007). Such scholars are suspicious that contemporary extensions of risk assessment and risk reduction will likely only reproduce, or may even exacerbate, the injustices of contemporary criminal justice policy under a more “objective” guise; and will fail even on their own terms at reducing crime (Harcourt, 2007).

Author Manuscript

There is of course a full spectrum of analysis, critique, and pragmatic positioning between these two perspectives. Much has been written on the use and misuse of risk assessment (e.g. Farabee et al., 2011; Hannah-Moffat, 2015; Hannah-Moffat et al., 2009; Monahan and Skeem, 2014, 2016; Skeem, 2013; Taxman and Marlowe, 2006; Taxman et al., 2006), and its social and political ramifications (e.g. Hannah-Moffat, 2004; Hannah-Moffat and Maurutto, 2010; Harcourt, 2015; O’Malley, 2012; Skeem and Lowenkamp, 2016). This article is not an attempt to summarize or supplant that rich dialogue. Rather, our aim is to excavate what we believe to be some of the buried conceptual roots of risk assessment’s contested terrain.

Author Manuscript

There is also an exhaustive empirical literature comprising many dozens of meta-analyses and systematic reviews that quantitatively evaluate the predictive validity and utility of various risk assessment instruments, often with conflicting conclusions (for a meta-review of these meta-analyses and systematic reviews see Singh and Fazel, 2010). The present article is not an attempt to comprehensively summarize the technical and empirical literatures. Rather, we draw on key insights from these literatures to illustrate the conceptual and theoretical issues on which this article focuses.

Author Manuscript

Central to our argument is that criminological risk assessment as a *way of knowing* is necessarily linked to *practices* based on this knowledge. Yet the ways in which this link is made and understood (both by its proponents and opponents) is often oversimplified (Hannah-Moffat, 2015; Hannah-Moffat et al., 2009), leading either to overconfidence or to poorly articulated opposition. The goal of this article, then, is to clarify the different dimensions of risk, and to articulate how “risk reduction” contains simultaneously the possibility of progressive reform and the perils of deepening inequality and ineffective intervention.

Author Manuscript

The stakes are not trivial, as the widespread acceptance and expansion of risk assessment in criminal justice policy is outpacing the theory and evidence to support it (Desmarais and Singh, 2013; Desmarais et al., 2016; Hannah-Moffat, 2012; Lowenkamp and Whetzel, 2009). The use of risk assessment is expanding from recidivism prediction to pre-trial processing, sentencing, and policing (Desmarais and Singh, 2013; Desmarais et al., 2016; Gottfredson and Moriarty, 2006; Lowenkamp and Whetzel, 2009; Storey et al., 2014; Summers and Willis, 2010; Trujillo and Ross, 2008; VanNostrand and Keebler, 2009). Yet the application of risk knowledge is often haphazard: jurisdictions frequently deploy pre-existing screening tools in settings for which they were neither designed nor calibrated (Lowenkamp et al., 2008), and legal and correctional professionals frequently do not understand the actuarial technologies upon which they base their decisions (Hannah-Moffat, 2012).

Before we begin, it is worth distinguishing “risk” as we discuss it here—as a shorthand for “actuarial thinking”—from two other conceptions of risk that are sometimes discussed in relation to criminal justice policy: “risk” as grand theory about our current socio-cultural condition, and “risk” as a narrow, technical concept. Most expansively, scholars have used risk to characterize society as a whole (Beck, 1992), and this conception of risk has been used within the criminal justice field as well. In *Governing through Crime*, for example, Jonathan Simon (2007) argues that the emergence of crime as a central social issue in the 1970s was intimately related to the end of the welfare state. When government could no longer guarantee a degree of material prosperity, he suggested, it turned to what it *could* provide: a sense of security. While important for understanding the socio-historical context in which “risk” has emerged as a dominant discourse, this conception of risk is not the focus of the present discussion.

At the opposite extreme, risk is used as a narrow technical concept to denote the empirical probability that an event will occur over a certain period of time. According to this more technical perspective, risk is distinct from uncertainty in that risk connotes knowledge (albeit incomplete), whereas uncertainty refers to a future event about which even the probability is unknown (Adams, 1995; Bernstein, 1996; Garland, 2003b; Knight, 1921; Rothman et al., 2008). From this vantage point, the increasing salience of risk assessment and risk management in criminal justice (and beyond) can be seen as an indication of growing trust in and use of knowledge generated about the probability of future criminality.

This more technical definition of risk informs our discussion, as it is a foundation upon which quantitative methods to determine “risks” and “risk factors” are built, but it is not our primary focus. Indeed, despite this clear technical definition, in practice “risk” is used to refer to many distinct quantitative concepts (Rothman et al., 2008). As a result, across many disciplines that rely on quantitative methods there has been rich debate about how risk should be measured, interpreted, and applied. Among these, epidemiology has wrestled perhaps more than any other over the quantification of risk and its translation into policy (see, for example, Lieberman, 1997; Link and Phelan, 1995; Rose, 1985; Rothman et al., 2008; Shmueli, 2010). We draw on these lessons in discussing actuarial thinking in criminology and criminal justice practice.

For expository purposes, in what follows, we engage heavily with the “Risk-Needs-Responsivity” framework articulated in *The Psychology of Criminal Conduct* (Andrews and Bonta, 2010). This framework, and its nearly 700-page elaboration, is representative of the dominant discourse of risk assessment and risk reduction. In turn, the Level of Services Inventory (LSI) and its various iterations (Andrews et al., 2004) is the risk instrument around which the Risk-Needs-Responsivity framework was constructed (Andrews et al., 2004). This instrument, while admittedly one of many, will stand in as the exemplar of the risk assessment framework in our analysis: it alone is employed by roughly 900 corrections agencies in North America (Lowenkamp et al., 2009) and shares much in common with the 19 other risk assessment instruments validated in US corrections settings (Desmarais and Singh, 2013; Desmarais et al., 2016). We choose to focus our analysis on what these risk assessment instruments share in common, rather than on their differences in content and application, because our concern is with the implications of the shared practice that their use

entails. We argue that the statistical models on which the theory and practice of risk assessment are based often conflate prediction and causation, and often hide population drivers of crime and criminal justice involvement.

In turn, we examine *Against Prediction* (Harcourt, 2007), a rigorous and influential critique of actuarialism in contemporary reform efforts by a widely respected theorist. Harcourt's argument highlights two additional problems with risk assessment—that risk instruments may *influence* the environments they purport to measure; and that the use of predictive techniques may distort our conceptions of just punishment. We argue that while each of these problems should be taken into consideration by scholars and practitioners, Harcourt's radical solution—abandoning risk thinking altogether—is likely equally undesirable. We conclude with our own perspective on how to discern between the policies and practices that analyze and respond to risk in appropriate ways and those circumstances in which “risk thinking” should be avoided.

The risky production of actuarial knowledge

A central premise of the Level of Services Inventory and other instruments like it is that they focus specifically on risk factors on which interventions might be focused, and so are (in theory) relevant not only to the *assessment* of risk but also to its *reduction* (Andrews and Bonta, 2010: 314–317). The Level of Services Inventory was developed based on research described by its originators as a “radical empirical approach to building theoretical understanding” (Andrews and Bonta, 2010: 132). Researchers identified variables that were most strongly correlated with re-arrest among individuals under community corrections supervision (Andrews and Bonta, 2010: 132–133), and then used those variables to categorize individuals into various risk groups for targeted treatment designed to reduce recidivism. Their research suggests that four risk factors consistently predict criminal conduct in almost any justice-involved sample: a history of antisocial behavior; antisocial personality pattern; antisocial cognition; and antisocial associates (Andrews and Bonta, 2010: 131). Versions of these risk factors appear in almost all of the 19 risk assessment instruments validated in US corrections settings (Desmarais et al., 2016). Furthermore, there is evidence that intervening upon criminogenic risk factors has some impact on criminal justice outcomes such as recidivism (e.g. Lowenkamp et al., 2006; Skeem et al., 2011). Because programs that ostensibly intervene upon criminogenic risk factors change target outcomes, these factors are treated—implicitly or explicitly—as causes of those outcomes.

An entire theoretical framework has emerged around these predictors, as suggested by the title of the authors' influential *The Psychology of Criminal Conduct* (Andrews and Bonta, 2010). Andrews and Bonta (2010: 133) admit that their approach might be confused with “dustbowl empiricism”, but argue that it is theoretically fruitful in that it “lead[s] to a deeper theoretical appreciation of criminal conduct” and is “practically useful in decreasing the human and social costs of crime” (2010: 133).

Without a doubt, the approach spearheaded by Andrews and Bonta has inspired new optimism within the field of criminal justice scholarship and practice that certain high-risk individuals can benefit from interventions focused on these “Big Four” risk factors (Cullen,

2011; Cullen and Gendreau, 2001; Gendreau et al., 2004; Lowenkamp et al., 2006). But there are two conceptual problems with these sorts of risk assessments. First, risk assessment instruments used to target individuals for treatment demand a *causal* theory of crime, yet the methodological foundations for these instruments are *predictive* rather than causal in nature. Second, risk assessment instruments arising from analysis of inter-individual variation in offending and arrest are necessarily blind to the drivers of *population* variation in offending and arrest, and so preclude any sorts of theorizing about or intervention on these population drivers.

Predictive statistics, causal inference

Andrews and Bonta claim to have identified an empirically driven, causal explanation for criminal conduct. This is, potentially, the most important strength of their approach. Previous risk assessment instruments, which did not make such causal claims, were of limited use to policy makers and practitioners. Such instruments might guide pre-trial detention decisions or be used to justify longer prison sentences for the sake of incapacitation, but they had little relationship to treatment or change at the individual level.

Nevertheless, throughout their work, Andrews and Bonta discuss risk assessment in ways that conflate prediction and causation. This slippage is common as well in practice, as noted by Hannah-Moffat (2012): despite the fact that risk instruments such as the LSI generate likelihoods based on group averages, criminal justice practitioners often interpret these outputs as administrative certainties. The distinction between causation and prediction can be a subtle one (Broadbent, 2011; Greenland, 2012; Shmueli, 2010), and is often obfuscated by the statistical modeling techniques underlying criminogenic risk assessment. But while statistical methods are inherently agnostic to predictive versus causal inferences, researchers who talk about manipulating or intervening upon risk factors in order to change outcomes cannot be agnostic; they are necessarily dealing in causation.

It is well established that a misspecified (“wrong”) quantitative causal model can have higher predictive validity than a model that correctly specifies the “true” causal relationship between independent and dependent variables (Shmueli, 2010).¹ Under certain assumptions, all causal effects will have some predictive value but not all predictive associations will have causal value. More often, the predictive value of a variable on an outcome partly reflects something causal, but may also reflect (perhaps much stronger) non-causal factors. In the criminal justice context, the fact that antisocial cognitions predict arrest, for example, does not help us understand how much of the association between antisocial cognitions and arrest is causal versus merely predictive.

Figures 1 and 2 present hypothetical models of the causal structure of select criminogenic risk factors and arrest, oversimplified for the sake of exposition. Figures 1 and 2 present

¹A heuristic example of a non-causal predictor is the association between carrying a lighter and lung cancer (Hernán and Robins, 2016). Carrying a lighter may predict lung cancer, but not because being around lighters causes lung cancer; rather, knowing that someone carries a lighter increases our knowledge about whether or not they smoke cigarettes. Cigarette smoking causes one to carry a lighter and also causes lung cancer: the lighter is only associated with (i.e. predictive of) lung cancer because they share a common cause. One might object and argue that banning lighters might reduce smoking, which might reduce lung cancer; however, this would represent an alternative causal model for lighters in which carrying a lighter causes smoking, rather than smoking causing one to carry a lighter.

these structures in the form of directed acyclic graphs (DAGs), in which solid, arrow-headed lines indicate causation (and its direction) and dotted lines represent non-causal association (Greenland et al., 1999; Shrier and Platt, 2008). In Figure 1a, antisocial behavior causes affiliation with antisocial peers (e.g. “birds of a feather”) and re-arrest. The association between antisocial peers and re-arrest might be large, but it represents a small causal effect. This occurs if the causal effect of antisocial behavior on peers and on re-arrest is very strong, in which case a measure of antisocial peers will pick up some of the effect of antisocial behavior (i.e. antisocial behavior is a strong confounder). In this scenario, it would be ineffective to intervene upon antisocial peers.

In Figure 1b, having antisocial peers causes antisocial behavior (e.g. by “enabling”), which causes re-arrest. In this model, peers still have some effect on re-arrest through mechanisms other than behavior. Here, it may make sense to intervene upon antisocial peers and block both the development of antisocial behavior (and thus re-arrest) in addition to the direct effect of antisocial peers on re-arrest. What is important about these first two examples is that while the associations between antisocial peers and re-arrest may be identical in both DAGs, they have very different implications for action.

Figure 2 illustrates a different scenario in which prediction and causation might be confused, this time by introducing a factor operating outside the individual (i.e. an aspect of social context) into the causal framework. Here, an intensive policing initiative results in more arrests, but also causes resentment and negative feelings about law enforcement in the target community. In this case, there is no causal relationship between antisocial cognitions and arrest, but there is a strong association. Intervening on antisocial cognitions, however, will have no effect on arrests, because these cognitions and higher arrest rates are merely effects of a common cause.

In the preceding discussion, we distinguished between prediction and causation, but provided examples where—regardless of the language used—practitioners’ interest should be in causal effects. Granted, there are some criminal justice contexts in which we might be interested primarily in predicting a future outcome. But manipulating risk factors to achieve intended outcomes requires causal knowledge of these relationships, and methods for estimating causal effects require different standards for validity than methods for actuarial prediction. When these two different purposes and methodological standards are confused or conflated, actuarial thinking can lead to problematic theorizing and ineffective practice.

Hiding population drivers of crime

A second problem with Andrews and Bonta’s framework is that it essentially hides drivers of crime (and exposure to the criminal justice system) at the population level. We discuss three ways in which this is manifested: (1) the studies on which the framework is based measure inter-individual variation among a selected population; (2) the framework concentrates on proximate causes of crime (and exposure to the criminal justice system) without considering the ways in which more distal causes put people at risk of proximate risks; and (3) the distinctions that Andrews and Bonta draw between “static” and “dynamic” risk factors presume a world in which population attributes are static and cannot be intervened upon. Although these issues have been acknowledged and discussed by

proponents of risk assessment and even the creators of the Risk-Needs-Responsivity framework, they have still led to inferential overreach (e.g. Andrews and Bonta, 2010: 40, 53, 131–138; Gutierrez et al., 2013).

First, the predictors on which Andrews and Bonta focus are only predictors of *inter-individual* variation in criminal activity. The authors discount sociological accounts of criminal conduct such as those that rely on theories about class, race, and neighborhood (Andrews and Bonta, 2010: 306–307). They argue that these variables do not predict variation in rates of individual recidivism as well as the criminogenic risk factors discussed above. Yet, as other scholars have noted (e.g. Hannah-Moffat, 2012), they take for granted that their studies almost exclusively compare individuals *within* a selected population—usually samples of individuals who are, or have at some point been, under some form of correctional supervision. This means that Andrews and Bonta conflate the subtle but important distinction between “causes of crime” and “causes of individual differences in crime”.

The distinction is important because the causes of individual differences in committing a crime or being arrested are almost certainly different from the causes of the incidence rate of crime (or exposure to the criminal justice system) across different populations or time periods. While the criminogenic risk patterns of individuals in a given area may predict who *among them* is more likely to engage in crime or be arrested (i.e. inter-individual variation), they almost certainly do not explain differences between different groups or areas (between-population variation) or over time. If the causes of these three types of variation were the same, we could understand differences in crime between populations and over time by documenting differences in the incidence and prevalence of individual-level risk factors such as antisocial personality pattern or antisocial cognitions. Mass incarceration as a historical phenomenon could be understood as merely the aggregation of changes in individual risk profiles. It seems much more likely, however, that the social and political drivers of mass incarceration are entirely (or almost entirely) separate from the factors that explain inter-individual variation in crime today.

The broader point is that the causes of a distribution are rarely the same as the causes of an individual's place within a distribution (Lieberson, 1997; Rose, 1985). As Lieberson (1997: 19) puts it, before one tries to understand individual outcomes, “[o]ne has to first consider any basic driving force that generates a certain outcome independent of the characteristics of the units in the population”. Thus, in order to understand differences in population distributions, Rose (1985) argues that we must study characteristics of populations, not characteristics of individuals. The trouble is that our statistical methods tend to be best suited for identifying causes of the latter, when we are often actually interested in causes of the former (Schwartz and Carpenter, 1999). This is because the “hardest cause to identify is the one that is universally present, for then it has no influence on the distribution of disease” (Rose, 1985: 33)—a phenomenon referred to as “ubiquitous exposure”.²

²For example, in a population of heavy smokers, lung cancer would appear to be a genetic disease (Rose, 1985).

For example, on average, socioeconomic status (SES) may not be a good predictor of individual arrest among a group of people with current or prior involvement in the criminal justice system, since a large majority of justice-involved people have low SES. Likewise, on average, SES may not be a good predictor of individual arrest among a group of people with virtually no connection to the justice system (since they would likely share similar levels of higher SES). But comparing these two groups would reveal that SES is a *strong* predictor of group differences in criminal justice system contact. Yet, for Andrews and Bonta (2010: 79, 93), it is a “myth” that the “roots of crime are buried deep in structured inequality”. They go on to cite the results of numerous meta-analyses, which find that lower-class origin is only weakly associated with criminality, relative to indicators of antisocial propensity (2010: 62–66). But if such distal risk factors (e.g. poverty, low education attainment) are nearly ubiquitous for individuals who become involved in the criminal justice system, we would not expect them to be predictive of inter-individual variation in re-arrest.³

Second, and closely related, Andrews and Bonta focus on proximate causes of crime and arrest without considering the impact of antecedent distal causes on those proximate causes. That is, even when distal risk factors vary enough to be detectable, proximate risk factors are prioritized. Andrews and Bonta’s perspective thus deemphasizes the ways in which a distal cause of crime and arrest may put people at risk of multiple proximate risks. Epidemiologists make a related distinction between “proximate” risk factors and “fundamental causes” of disease—things like socioeconomic status or racial discrimination, which put people “at risk of risks” (Link and Phelan, 1995). Fundamental causes influence multiple diseases through multiple mechanisms, and are often treated as static background characteristics. Instead, factors that are “closer” to the outcome are understood as “causes”, while the causes of these proximate causes are left unexamined. Link and Phelan (1995: 85), critiquing modern epidemiology’s focus on individual risk factors such as health behaviors, argue, “efforts to reduce risk by changing behavior may be hopelessly ineffective if there is no clear understanding of the process that leads to [those behaviors]”. Said differently, interventions on a dynamic risk factor close to the outcome of interest will be short-lived if we do nothing to intervene on the causes of the risk factor. For example, rates of obesity and diabetes are high among low-income communities; however, exhorting poor people to eat non-processed foods without addressing the limited access to such foods in their communities is unlikely to have an impact (Link and Phelan, 1995: 85).

This is the case even though distal causes, further away from an outcome on a causal pathway, are likely to be only weakly associated with the outcome once more proximate factors are included in the model. This basic concept is simple to understand in nontechnical terms: distal factors take many more intermediate steps to arrive at an outcome, and these intermediate steps each “absorb” some of the statistical association on the way to the outcome. In contrast, proximate factors require far fewer intermediate steps, and less of the

³It is interesting to note that in some passages, Andrews and Bonta appear to understand this, particularly when they explain how political economy, social structure, and culture fit into their general personality and social psychological theory of crime. For example, regarding the aforementioned factors, they recognize that “[b]ecause they are constants, they are distal background contextual conditions that cannot account for variation in individual conduct within particular social arrangements” (2010: 137). Yet they seem to set this understanding aside when critiquing sociological or population approaches to criminological theory and methodology.

association is absorbed along the way. Thus, the findings from the meta-analysis above are exactly what we would hypothesize in a fully elaborated causal model of criminality.

Finally, because Andrews and Bonta are interested in explaining (and intervening on) inter-individual variation rather than population-level variation, a similar individual-level bias is embedded in the distinction that they make between static and dynamic risk factors (e.g. Andrews and Bonta, 2010: 27–30). Most simply, this is a distinction between factors that cannot be changed (i.e. criminal histories, or the physical attributes used to racialize people) versus factors that can be manipulated (i.e. our job skills, cognitions, or social networks).

The problem here is not so much with the distinction itself as with the assumptions about where the boundary lies between static and dynamic factors. Much of the risk literature today, focused on individual intervention, views questions about certain “static” factors as beyond its scope. Yet some characteristics of neighborhoods and communities that seem “static” at the level of any individual may be dynamic at another level, changeable through policy or social action. Maintaining a rigid distinction between static and dynamic factors means that we might fail to critically conceptualize or question some of the social categories we are interested in measuring and ultimately changing (Schwartz et al., 2011). For example, when we talk about “race” in any sophisticated sense, we are not referring merely to a person’s skin color and certainly not to any biological given, but rather to a person’s lived experience as the occupant of an imposed, racialized social position. The way that individuals are treated by other people and institutions is of course manipulable. The point is that labeling a factor as “static” is very often a political calculus, not a reflection of an empirical reality.

Beyond biasing our understanding of the “causes” of crime, the prioritization of inter-individual variation, proximate causes, and individual-level dynamic risk factors privileges individual-based interventions rather than population-based or structural interventions. We focus on intervening on those individuals at highest risk (that is, on removing the right tail of the distribution), rather than intervening on an entire population in order to shift the mean of the entire distribution (Rose, 1985). We are absolved from intervening on factors that put people at risk of individual-level risks—things like housing, unemployment, economic inequality, community disinvestment, and racial discrimination, for instance—even as they shape distributions of crime and justice system contact.

Should we be *against prediction*?

The problems with risk assessment that we identify above—its tendency to conflate prediction and causation, and its tendency to focus on inter-individual variation rather than population-level variation or distal causes—should give pause to those scholars and practitioners who believe they have found in instruments such as the LSI a silver bullet to correctional classification and intervention. And yet each critique might be incorporated into future generations of risk assessment and risk reduction: scholars and practitioners might work more carefully to distinguish between predictive factors and causal factors, and to create models of assessment and intervention that include causes of population variation as well as individual variation and target distal as well as proximate risk factors.

Another set of critiques, elaborated perhaps most extensively by Bernard Harcourt (2007) in his book, *Against Prediction*, strike more deeply at the very foundations of risk assessment in criminal justice, although his argument is built on a narrow set of hypothetical claims about police profiling. First, he argues, risk assessment might actually increase crime across a population, if moving dollars from controlling one group to another group increases crime among the lower-risk group more than it reduces crime among the higher-risk group (Harcourt, 2007: 123). Second, he argues, risk assessment has the potential to “ratchet-up” social disparities within the criminal justice system, even if it reduces offending at a population level, when “law enforcement relies on the evidence of correctional traces—arrests or convictions—in order to reallocate future law enforcement resources” (Harcourt, 2007: 156). Finally, and most generally, Harcourt suggests that the use of predictive techniques distorts our conceptions of just punishment, as we move away from punishment based on what someone has done and toward punishment based on what someone is expected to do (Harcourt, 2007: 3).

Risk assessments and the worlds they create

Harcourt’s first two concerns are centrally about the reflexivity of risk. Risk assessments are not only descriptive but constitutive of the social world (Fourcade and Healy, 2013). This broad point is consistent with recent theorizing about “reflexive modernization” (Beck et al., 1994; but see Alexander, 1996), which explores the ways in which our scientific knowledge of the world constrains and enables action within it. Each of Harcourt’s concerns is valid theoretically, but Harcourt’s examples do not hold up to the existing empirical evidence. This matters because Harcourt connects these hypothetical situations to a sweeping indictment of the use of prediction in criminal justice as a whole. He concludes that while progress in the fields of “communication, transportation, and medicine are models of success, and the advances we have made in those areas are simply remarkable [...] the study of our political, social, and economic organization [...] [has] not produced such results” (2007: 239). Harcourt (2007: 239) argues instead that we should “quell the desire to put these prediction instruments to use” in criminal justice policy, “the most devastating area of social life”. Rather than wrestle with the difficulties inherent in moving from predictive statistics, to causal inference, to intervention, he suggests that data have no place in criminal justice practice. Harcourt may indeed be right that contemporary applications of risk assessment have had “devastating” consequences in criminal justice, but these consequences are a result of the ways in which they have been applied, not of the methods themselves. Harcourt’s totalizing critiques of contemporary risk assessment are ultimately as unjustifiable as the rabid enthusiasm of its strongest proponents.

In Harcourt’s first scenario, a lower-risk group, classified as such and policed less closely, might begin offending more—and become higher risk. In turn, a higher-risk group, classified as such and policed more closely, might begin offending less—and become lower risk. This is an important point, and one that has been recognized by others in the literature. As Garland (2003b: 54) notes, many of the risks we hope to reduce “are altered as soon as we identify them as such”. Indeed, in the context of epidemiology, we often hope that the identification of risk factors will lead reflexively to changes in behavior or changes in

resource allocation (Rose, 1985). In other words, the reflexivity of risk is in many ways the *motivation* for risk-based models in the health sciences.

Harcourt continues that, depending on the “relative elasticities” of each group to changes in policing, total crime might actually increase. This is mathematically true, as Harcourt goes to great lengths to demonstrate. But he provides no evidence that predictive policing actually *has* led to increases in crime rates. Indeed, several decades of policing research provide compelling evidence that people do not adjust activity based on even rather large increases in random patrols; much evidence suggests instead that low-level, evenly distributed policing across a population has almost no impact on anyone (Skogan and Frydl, 2004). Rather, in order for policing to have *any impact at all* on anyone, it must be highly concentrated (Braga and Weisburd, 2010; Kelling et al., 1974; Kleiman, 2009; Weisburd and Eck, 2004). In his terms, the elasticity of crime rates to policing resources is essentially zero at all but the highest levels of police resource concentration. As a result, diverting police resources from low-risk groups toward high-risk groups likely has almost no marginal impact among the low-risk populations. Focusing police resources on risky people and places means the difference between doing nothing (for all) and doing something (for some).

Harcourt’s second concern, based on a separate set of logical assumptions (now crime is assumed to be *inelastic* to changes in policing), is that concentrating police resources on riskier groups will lead to a “ratchet effect” in which those riskier groups will become more and more disproportionately punished. Under this scenario, risk assessments are skewed by existing distributions of the carceral population, which in turn lead to *further* disproportionality in carceral distributions, which lead to further disproportionalities in risk assessments, and so on. This, again, is a valid concern theoretically. For example, if conviction rates are used to generate risk assessment measures, but racism pervades every step of the criminal justice process, then these assessments will inevitably overstate the salience of race in criminal risk. Depending on how these assessments are then used, they might deepen the racism of the justice system by increasing the targeting of supervision (or other sorts of sanctions) on racialized groups.

Again, however, the existing empirical evidence does not support Harcourt’s legitimate hypothetical concern. Historical evidence to support a racialized “ratchet effect” in criminal justice, the most salient example of a potential ratchet (and the one on which Harcourt grounds his argument), is ambiguous. If the ratchet were a straightforward phenomenon, we would expect to see consistent growth in the *ratio* of minority to white incarceration rates over time. This does not appear to be the case. Muller (2012) has shown that the most consistent and dramatic growth in the ratio of minority to white incarceration rates in the United States took place between 1880 and 1950, long before the rise of mass incarceration. In more recent years, growth in this disproportionality seems to have been either uneven or nonexistent. Some scholars argue that racial disparities did not grow at all between 1980 and 2000 (Western, 2006: 30), whereas others have argued that such disparities rose in the 1970s and 1980s, peaked in the 1990s, and have been slowly declining since (Oliver, 2001, 2008). Whatever the case, if the ratchet phenomenon is to be a useful concept, a more nuanced, trend-responsive explanation for this “fluctuating ratchet” would be required.

There are other analytic problems with Harcourt's "ratchet effect" as well. In Harcourt's hypothetical situation, he assumes that a majority group offends at a rate of 6 percent and a minority group offends at a rate of 8 percent. Harcourt points out that, if police are searching people within each group "randomly" and trying to maximize efficiency (i.e. hoping to maximize "hit rate"), they will tend toward policing the second group more frequently. Harcourt actually *underestimates* the ratchet effect under the assumptions of his own model (Margalioth, 2008). Under his assumptions, efficient police (interested in maximizing their hit rate) should search the second group *exclusively* from the beginning—they will always have higher success (8 percent) with the minority group than with the majority group (6 percent). Yet this model contains another questionable assumption—namely, that police typically search people *entirely randomly* within any group. A more realistic scenario, at least most of the time, is that police use several different group-level differences or characteristics—or, worse, different stereotypes and biases—to raise or lower their level of suspicion concerning members of particular groups (based on age, race, gender, clothing, associates, neighborhood, patterns of behavior, criminal history, car color, etc.).

Under Harcourt's model of random policing, "hit rates" by group (or characteristic)—the rates at which police search someone "successfully", finding evidence of crime—stay the same no matter how searches are allocated across groups. Yet this is almost certainly incorrect. Rather, "hit rates" are likely elastic to the rates that a police officer *searches* among a certain group. For example, as an officer's attention (or suspicion) increases with regard to a particular group (or characteristic), the officer will be marginally less discriminating about whom to search within that group, and the hit rate within that group will decline. Conversely, as an officer's attention (or suspicion) declines with regard to a group, the officer will be marginally more discriminating about whom to search within that group, and the hit rate will rise. Under Harcourt's model, "hit rates" between groups will never equalize, and so the police will wind up "fishing"—to borrow Harcourt's own metaphor—exclusively from the higher-crime group (Harcourt, 2007: 47–48). In contrast, once we recognize that there is elasticity in the relationship of hit rates to search rates, we can appreciate that it would not be in any officer's interest to patrol one group *exclusively*. Rather, the officer would be efficiency-maximizing when hit rates are the same, and this would occur at some middle-range distribution of searches between groups. Granted, even under these more realistic conditions, there will still be disproportionality between underlying rates of criminality and their representation in stops and arrests, depending on the relative elasticities of hit rates to search rates. However, this should be an *immediate* disproportionality that does not necessarily get worse over time, provided that hit rates are the measure of police efficiency.⁴

The problem of justice

Even if Harcourt is wrong about a ratchet effect, he is correct that "efficient" policing will likely exacerbate group-level inequality in the criminal justice system. His critique of actuarialism now transitions from one based on the inefficiency of risk assessment to one

⁴Harcourt bases his own "ratchet" on the idea that police consistently adjust their understandings of risk based on the current distribution of individuals under correctional supervision. This does not seem consistent with either his own theory or with any empirical evidence.

based on its injustice. If different groups offend at different rates, and the criminal justice system takes these differences into consideration in its interventions, then the population on which it intervenes will include a disproportionate share of the higher-offending group (relative to differences in offending rates). Furthermore, Harcourt (2007: 162–163) is certainly correct that, in the context of *race*—the context that Harcourt emphasizes throughout his book—such a ratchet “contributes to an exaggerated general perception in the public imagination and among police officers of an association between being African American and being a criminal”.

Harcourt (2007: 237) writes that the ratchet effect violates “a core intuition of just punishment—the idea that anyone who is committing the same crime should face the same likelihood of being punished regardless of their race, sex, class, wealth, social status, or other irrelevant categories”. Randomness in policing, Harcourt (2007: 237–238) suggests, is the clearest way to satisfy this core intuition: “[t]he only way to achieve our ideal of criminal justice is to avoid actuarial methods and to police and punish color-blind, gender-blind, or class-blind [...] Randomization in this context is a form of random sampling [...]” This, again, is an important theoretical point. Minimizing risk is only one of many competing values with which people approach social problems in general, and criminal justice in particular (Douglas and Wildavsky, 1982; Renn, 1998), and there are circumstances in which it is undoubtedly considered less important than a competing value. As just one example, those convicted of murder have the lowest recidivism rates among any category of released offenders (Langan and Levin, 2002). Yet releasing these offenders earlier than others because of this small risk would (for many) violate a sense of proportionality, a sense of justice.

Yet four things should be said about the unjust disproportionality Harcourt thinks “randomness” will resolve. The first is that our feelings about it might be different if it were based on something *other* than a “risk factor” like race—a proxy for historical and ongoing systems of racial hierarchy, domination, and exploitation. That is, we might be more willing to have an “efficiency-maximizing” criminal justice system that had a disproportionate number of repeat offenders (or violent offenders, or white collar criminals) relative to all people engaged in behavior that would be considered criminal if discovered by police.⁵ The justifications for such disproportions would surely be qualitative, and socially derived, rather than quantitative or empirical.

Second, we would likely feel differently if the *intervention* being disproportionately allocated was not so exclusively punitive—for example, if the intervention were not arrest and conviction but access to housing, cognitive-behavioral therapy, and job training. To draw again on our analogies from epidemiology and public health, if an efficient distribution of heart surgeries meant that those at greatest risk of heart attack were disproportionately offered the service (to the detriment of those with lower, but still real, risk of heart attack), our evaluation of the allocation would almost certainly be different (though we still might object on different grounds).

⁵Harcourt (2007: 164–165) suggests that the ratchet effect might be just as damaging for repeat offenders—what he calls “recidivist criminality”.

Third, hit rates may indeed be a problematic measure of justice and efficiency in policing, but we must specify *why* this is the case, and specify reasonable alternatives. The disproportionate representation of certain groups within the correctional system (to which Harcourt points) is one potential injustice; the fact that maximizing hit rates may not minimize crime rates is a potential inefficiency. Perhaps as problematic from a justice perspective, however, is the high number of false positive stops that police may undertake within certain groups as a result of probabilistic stops, even if hit rates across groups are equal.

Even in the absence of the reflexive mechanisms Harcourt discusses, risk instruments have important effects through the different sorts of *wrong* predictions they inevitably produce. One important practical dimension regarding any risk assessment is its accuracy, and the consequences of inaccuracy. Colloquially, it may be tempting to view accuracy in terms of “percent of correct predictions”, but quantifying accuracy is much more complex. To illustrate, consider a predictive instrument that is 99.9 percent “correct”. If one predicts that 0 out of 10,000 people will commit a murder, but in truth 10 people out of 10,000 commit a murder, one would still be “99.9 percent correct”, because 9,990 out of 10,000 did not commit a murder. To overcome this problem, more sophisticated measures of predictive validity have been developed, such as sensitivity, specificity, positive predictive value, and negative predictive value.⁶

The preferred quantification of predictive validity depends on the question of interest, and the relative stakes of false positives and false negatives. During airport screenings we may be more tolerant of assessments with low specificity (high false positives) than we are during street stops, for example. For a deadly contagious disease, we may be particularly worried about missing anyone with the condition (i.e. false negatives), and want to make sensitivity as high as possible. Predictive validity is thus an exercise in tradeoffs, because as false positives go down, false negatives tend to go up, and vice versa. The way that these tradeoffs are made is again a political question, outside the scope of the assessment itself. But it is telling that these tradeoffs are not acknowledged nor discussed within much of the literature on risk, whether by proponents or critics. One thoughtful exception is the work of Fazel and colleagues (2012), which compares risk assessments for violence and antisocial behaviors to other common medical risk assessment tools, and argues that any such comparisons must take into account differences in economic, social, and civil rights implications between these domains.

That is, even if police are not being racially biased in their stops according to “hit rate”, we might object based on the intrusion that searches impose on the liberty of non-criminal group members.⁷ The irony, however, is that certain focused-deterrence strategies may *mitigate* the problem of false positives while *worsening* the problem of disproportionality, if

⁶Sensitivity is the proportion of people who test positive out of everyone who is truly positive (i.e. true positives divided by true positives and false negatives) and specificity is the proportion of people who test negative out of everyone who is truly negative (true negatives divided by true negatives and false positives). Positive predictive value, or precision, is the proportion of people who are truly positive given that they tested positive (true positives divided by true positives and false positives), and negative predictive value is the proportion of people who are truly negative given that they tested negative (true negatives divided by true negatives and false negatives).

⁷Harcourt (2007: 169–170) discusses this other consequence only briefly.

rates of offending are truly disparate. In other words, assuming no change in offending patterns, the more specific we can be about the groups at highest risk of crime, the *more disproportionate* the subsequent carceral population will be, though our “hit rates” will likely be higher, meaning fewer false positives.

Finally, Harcourt’s model of just punishment—that is, anyone committing the same crime facing the same likelihood of capture—depends upon a conception of random policing that supposes no elasticity of search rate to hit rate. In the policing context, he puts forward a “radical idea [...] to draw social security numbers by lottery and then have a full investigation of the person’s life” (2007: 238). In this model, police are assumed to be blind to *all* difference (the person actively robbing someone else is equally likely to be investigated as the person who is catatonic in a nursing home facility). It is true that, under these constraints, the population under the supervision of the criminal justice system would reflect proportionately the population of people involved in crime. But policing, under these circumstances, would likely serve *no other purpose than generating this proportionality*.

Most policing scholars make an assumption more in line with what we outline above—all else being equal, group-level hit rates will be elastic to group-level search rates. From this perspective, allocating policing resources equally across a population (a more realistic “randomness”) will lead to *different* likelihoods of being punished for the same crime, given different group-level crime (or criminalization) rates. In his *When Brute Force Fails*, for example, Kleiman (2009: 23) argues,

Since enforcement and prosecution resources are much more equally distributed than is crime, an offender who commits a crime where crime is common is less vulnerable to arrest, vigorous prosecution, and a stiff sentence than an offender who commits a crime in a more law-abiding neighborhood.

As a result, he suggests, such an equal enforcement strategy violates

the constitutional mandate of “equal protection under the laws,” if “equal protection” means that a crime against a poor or black person will be investigated as diligently, prosecuted as forcefully, and punished as severely as the same crime against a rich or white person.

(2009: 23)

The irony, he concludes, is that “[p]roviding something closer to actual equal protection of the laws would make the problem of disproportion in punishment worse, not better, unless and until higher per-crime punishment risks caused African American crime rates to fall” (2009: 24). In other words, if policing does not take into consideration any group or geographic differences in crime, it is actually unlikely to lead to the sort of justice to which Harcourt aspires.

When risk is not reductionist

There is a long history to the idea that “nothing works” in the field of criminal justice that is beyond the scope of this analysis. The more relevant question here is whether the skepticism epitomized by Harcourt is warranted. Just as the proponents of risk assessment have, by and

large, neglected the distinctions we draw above, so too have the most prominent critics of risk assessment neglected careful consideration of exactly what risk assessment is, what questions it can legitimately and appropriately answer, and how to implement “risk knowledge” when it is addressed to appropriate questions.

Rather than embrace actuarial thinking uncritically, or reject it entirely, our analysis recommends a deepening engagement with the assumptions underlying risk assessment. As scholars and practitioners we ought to be more discerning about those conditions under which risk might enhance public safety and cost-efficiency, and—in turn—about those circumstances in which risk thinking should be set aside for the sake of other values and priorities. This is a point made nicely by Pat O’Malley (2008: 453), who argues that “risk as an abstract technology is always shaped and given effect by specific social and political rationales and environments” and that “there is no obvious reason why risk cannot be inclusive and reformist rather than exclusionary and merely incapacitating”. In this concluding section we put forward explicitly what we have implied throughout this article, and outline guidelines or principles for producing and applying risk knowledge in criminal justice.

First, the production and use of risk knowledge seems most promising when used in the context of causal inference, so that it has relevance for harm reduction and treatment rather than only for surveillance and punishment. The focus of risk knowledge should thus be on causal, dynamic factors (and by dynamic we include community-level factors that contemporary instruments consider static) or on causal static factors that enhance our understanding of institutional and structural processes, rather than on individuals’ circumstances or characteristics. Granted, there are circumstances in which the use of risk knowledge for actuarial prediction may be appropriate. For example, decisions about pre-trial detention may be a case in which actuarial prediction *is* the most suitable aim of risk assessment. Under these circumstances, however, the predictive purpose should be justified explicitly—explanatory (causal) frameworks should not be imposed post hoc.

Second, the production and use of risk knowledge should be explicit about the context in which it was generated, situating risk factors in broader causal frameworks, and taking into consideration both individual-level and population-level differences and proximate and more distal risk factors. Ideally, interventions at a proximate-level (e.g. antisocial cognitions) might be strategically connected to interventions at a more fundamental level (e.g. community mistrust of the police). This is the insight behind interventions such as David Kennedy’s (2011) Ceasefire model, which link focused-deterrence policing strategies (a proximate-level intervention) with broader efforts to address histories of racial injustice and oppression (a structural intervention). Relatedly, those who use and apply risk knowledge should be conscious about the ways in which it is inevitably qualitative and value-laden (i.e. what is considered a risk, what is considered static or dynamic, whether interventions should be focused on individuals or populations, the relationship between the sensitivity and specificity of any assessment). They should be careful about the extent to which this risk knowledge perpetuates or reproduces existing inequality, either tacitly or overtly.

Finally, any application of risk knowledge must acknowledge that there are arenas of criminal justice policy in which predicting (and reducing) future harms should be considered less important than other social values. There is a very long history of scholarship on the proper functions of criminal justice sanction. But most would agree, at least in theory, that criminal justice systems have both expressive and utilitarian purposes. By designating some things as criminal, for example, justice systems highlight shared social values.

Considerations of justice are unavoidable in contemporary criminal justice as well.

Critics such as Harcourt are wise to emphasize that the production and application of risk knowledge should *serve*, rather than constitute, our conceptions of justice. Using risk knowledge to make existing systems run more efficiently should not be an end in itself, as it sometimes becomes (Feeley and Simon, 1992); rather, the production and application of risk knowledge should be used to understand and eliminate the inequities and injustices that the current penology perpetuates. On the other hand, Harcourt's *generalized* skepticism toward risk thinking seems unwarranted. In contrast to Harcourt, we believe that the social sciences can play an important role in criminal justice reform, even if they have not always done so in the past. And so we hope that current and future penologies incorporate the pragmatic and action-oriented use of responsible risk thinking, although we encourage researchers and policy makers to be more circumspect about the limitations (and indeed dangers) of this approach, and more open to population-based and structural interventions on the fundamental causes of crime.

Acknowledgments

This article has benefited tremendously from the advice and comments of Marshall Clement, Bruce Link, Sharon Schwartz, Chris Muller, and Sandra Smith.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/ or publication of this article: Reich thanks the Robert Wood Johnson Health & Society Fellowship Program for support.

Biographies

Seth J Prins is a postdoctoral research fellow in Sociomedical Sciences at Columbia University. His research interests include the criminalization of mental illness, the medicalization of crime, and the political-economic determinants of mood disorders. He received his PhD in Epidemiology at Columbia University and his MPH in Sociomedical Sciences at the Columbia University Mailman School of Public Health. His previous research appears in journals such as *Law and Human Behavior*, *Criminal Justice and Behavior*, *Psychiatric Services*, *Sociology of Health & Illness*, *Social Science & Medicine*, *Social Theory & Health*, *International Journal of Health Services*, *Addiction*, and *Community Mental Health Journal*.

Adam Reich is an assistant professor of sociology at Columbia University. He received his PhD in sociology from UC Berkeley in 2012, and was a Robert Wood Johnson Health & Society Scholar at Columbia from 2012 to 2014. His focus, broadly, is on economic and cultural sociology. Reich is the author of three books, including *Hidden Truth: Young Men*

Navigating Lives In and Out of Juvenile Prison (University of California Press, 2010). He is also the author of several peer-reviewed articles, which have appeared in journals such as the *American Journal of Sociology* and *Social Science & Medicine*.

References

- Adams J. Risk London: Routledge; 1995
- Alexander JC. Critical reflections on “reflexive modernization”. *Theory, Culture & Society*. 1996; 13(4):133–138.
- Andrews DA, Bonta J. *The Psychology of Criminal Conduct* Albany, NY: LexisNexis, Anderson Publishing; 2010
- Andrews DA, Bonta J, Wormith JS. *The Level of Service/Case Management Inventory (LS/CMI)* Toronto: Multi-Health Systems; 2004
- Beck U. *The Risk Society: Towards a New Modernity* Ritter M, translator London: SAGE; 1992
- Beck U, Giddens A, Lash S. *Reflexive Modernization: Politics, Tradition, and Aesthetics in the Modern Social Order* Stanford, CA: Stanford University Press; 1994
- Bernstein PL. *Against the Gods: The Remarkable Story of Risk* New York: Wiley; 1996
- Braga A, Weisburd D. *Policing Problem Places: Crime Hot Spots and Effective Prevention* Oxford: Oxford University Press; 2010
- Broadbent A. What could possibly go wrong? A heuristic for predicting population health outcomes of interventions. *Preventive Medicine*. 2011; 53(4–5):256–259. [PubMed: 21967964]
- Cullen FT. Beyond adolescence-limited criminology: Choosing our future. *Criminology*. 2011; 49(2): 287–330.
- Cullen FT, Gendreau P. From nothing works to what works: Changing professional ideology in the 21st century. *The Prison Journal*. 2001; 81(3):313–338.
- Desmarais SL, Singh JP. *Risk Assessment Instruments Validated and Implemented in Correctional Settings in the United States* New York: Council of State Governments Justice Center; 2013
- Desmarais SL, Johnson KL, Singh JP. Performance of recidivism risk assessment instruments in U.S. correctional settings. *Psychological Services*. 2016; 13(3):206–222. [PubMed: 27267819]
- Douglas M, Wildavsky A. *Risk and Culture: An Essay on the Selection of Technological and Environmental Dangers* Berkeley, CA: University of California Press; 1982
- Farabee D, Zhang S, Yang J. A preliminary examination of offender needs assessment: Are all those questions really necessary? *Journal of Psychoactive Drugs, SARC Suppl*. 2011 Sep.;51–57.
- Fazel S, Singh JP, Doll H, et al. Use of risk assessment instruments to predict violence and antisocial behaviour in 73 samples involving 24,827 people: Systematic review and meta-analysis. *BMJ: British Medical Journal*. 2012; 345(7868):1–12.
- Feeley MM, Simon J. The new penology: Notes on the emerging strategy of corrections and its implications. *Criminology*. 1992; 30(4):449–474.
- Fourcade M, Healy K. Classification situations: Life-chances in the neoliberal era. *Accounting, Organizations and Society*. 2013; 38(8):559–572.
- Garland D. Penal modernism and postmodernism. In: Blomberg TG, Cohen S, editors *Punishment and Social Control* New York: Transaction Publishers; 2003a
- Garland D. The rise of risk. In: Ericson RV, Doyle A, editors *Risk and Morality* Toronto: University of Toronto Press; 2003b
- Gendreau P, French SA, Gionet A. What works (what doesn’t work): The principles of effective correctional treatment. *Journal of Community Corrections*. 2004; 13(3):4–6.
- Gottfredson SD, Moriarty LJ. Statistical risk assessment: Old problems and new applications. *Crime and Delinquency*. 2006; 52(1):178–200.
- Greenland S. Causal inference as a prediction problem: Assumptions, identification and evidence synthesis. In: Bernardinelli L, Dawid P, Berzuini C, editors *Causality: Statistical Perspectives and Applications* Chester: John Wiley & Sons; 2012

- Greenland S, Pearl J, Robins JM. Causal diagrams for epidemiologic research. *Epidemiology*. 1999; 10(1):37–48. [PubMed: 9888278]
- Gutierrez L, Wilson HA, Rugge T, et al. The prediction of recidivism with Aboriginal offenders: A theoretically informed meta-analysis. *Canadian Journal of Criminology and Criminal Justice*. 2013; 55(1):55–99.
- Hannah-Moffat K. V. Gendering risk at what cost: Negotiations of gender and risk in Canadian women's prisons. *Feminism and Psychology*. 2004; 14(2):243–249.
- Hannah-Moffat K. Actuarial sentencing: An unsettled proposition. *Justice Quarterly*. 2012; 30(2):270–296.
- Hannah-Moffat K. The uncertainties of risk assessment: Partiality, transparency, and just decisions. *Federal Sentencing Reporter*. 2015; 27(4):244–247.
- Hannah-Moffat K, Maurutto P. Re-contextualizing pre-sentence reports: Risk and race. *Punishment & Society*. 2010; 12(3):262–286.
- Hannah-Moffat K, Maurutto P, Turnbull S. Negotiated risk: Actuarial illusions and discretion in probation. *Canadian Journal of Law and Society/Revue Canadienne Droit et Societe*. 2009; 24(3): 391–409.
- Harcourt BE. *Against Prediction: Profiling, Policing, and Punishing in an Actuarial Age* Chicago, IL: University of Chicago Press; 2007
- Harcourt BE. Risk as a proxy for race: The dangers of risk assessment. *Federal Sentencing Reporter*. 2015; 27(4):237–243.
- Hernán MA, , Robins J. *Causal Inference* Boca Raton, FL: Chapman and Hall/CRC; 2016
- Kelling GL, , Pate T, , Dieckman D. , et al. *The Kansas City preventive patrol experiment: A summary report* Washington, DC: The Police Foundation; 1974
- Kennedy D. *Don't Shoot: One Man, a Street Fellowship, and the End of Violence in Inner-City America* New York: Bloomsbury Publishing; 2011
- Kleiman MAR. *When Brute Force Fails: How to Have Less Crime and Less Punishment* Princeton, NJ: Princeton University Press; 2009
- Knight FH. *Risk, Uncertainty, and Profit* Boston, MA: Houghton Mifflin; 1921
- Langan PA, , Levin DJ. *Recidivism of prisoners released in 1994* Washington, DC: Bureau of Justice Statistics Office of Justice Programs, US Department of Justice; 2002 Available at: www.jstor.org/stable/10.1525/fsr.2002.15.1.58
- Lieberson S. Modeling social processes: Some lessons from sports. *Sociological Forum*. 1997; 12(1): 11–35.
- Link BG, Phelan J. Social conditions as fundamental causes of disease. *Journal of Health and Social Behavior*. 1995:80–94. Extra Issue. [PubMed: 7560851]
- Lowenkamp CT, Whetzel J. The development of an Actuarial Risk Assessment Instrument for U.S. pretrial services. *Federal Probation*. 2009; 73(2):33–36.
- Lowenkamp CT, Latessa EJ, Smith P. Does correctional program quality really matter? The impact of adhering to the principles of effective intervention. *Criminology Public Policy*. 2006; 5(3):575–594.
- Lowenkamp CT, Lemke R, Latessa E. The development and validation of a pretrial screening tool. *Federal Probation*. 2008; 72(3):2–9.
- Lowenkamp CT, Lovins B, Latessa EJ. Validating the Level of Service Inventory— Revised and the Level of Service Inventory: Screening version with a sample of probationers. *Prison Journal*. 2009; 89(2):192–204.
- Margalioth Y. Looking at prediction from an economics perspective: A response to Harcourt's *Against Prediction*. *Law and Social Inquiry*. 2008; 33(1):243–252.
- Monahan J, Skeem JL. Risk redux: The resurgence of risk assessment in criminal sanctioning. *Federal Sentencing Reporter*. 2014; 26(3):158–166.
- Monahan J, , Skeem JL. *Risk assessment in criminal sentencing* 2016 Nov. 125 Available at:
- Muller C. Northward migration and the rise of racial disparity in American incarceration, 1880–1950. *American Journal of Sociology*. 2012; 118(2):281–326.
- O'Malley P. Experiments in risk and justice. *Theoretical Criminology*. 2008; 12(4):451–470.

- O'Malley P. Crime and Risk [online edn] London: SAGE; 2012 Risk and crime control; 2352 Available at:
- Oliver PE. Racial disparities in imprisonment: Some basic information. *Focus*. 2001; 21(3):28–31.
- Oliver PE. Racial patterns in state trends in prison admissions 1983–2003. Drug and non-drug sentences and revocations 2008 Available at: www.ssc.wisc.edu/~oliver/RACIAL/StateTrends/RacialPatterns_Intro_National.pdf
- Renn O. Three decades of risk research. *Journal of Risk Research*. 1998; 1(1):49–71.
- Rose G. Sick individuals and sick populations. *International Journal of Epidemiology*. 1985; 14(1):32–38. [PubMed: 3872850]
- Rothman KJ, Greenland S, Lash TL. *Modern Epidemiology* Philadelphia, PA: Lippincott, Williams and Wilkins; 2008
- Schwartz S, Carpenter KM. The right answer for the wrong question: Consequences of type III error for public health research. *American Journal of Public Health*. 1999; 89(8):1175–1180. [PubMed: 10432902]
- Schwartz S, Gatto NM, Campbell UB. What would have been is not what would be: Counterfactuals of the past and potential outcomes of the future. In: Ornstein K, Keyes KM, Shrout PE, editors *Causality and Psychopathology: Finding the Determinants of Disorders and Their Cures* New York: Oxford University Press; 2011
- Sherman LW. The power few: Experimental criminology and the reduction of harm. *Journal of Experimental Criminology*. 2007; 3(4):299–321.
- Shmueli G. To explain or to predict? *Statistical Science*. 2010; 25(3):289–310.
- Shrier I, Platt RW. Reducing bias through directed acyclic graphs. *BMC Medical Research Methodology*. 2008; 8:70. [PubMed: 18973665]
- Simon J. *Governing through Crime: How the War on Crime Transformed American Democracy and Created a Culture of Fear* New York: Oxford University Press; 2007
- Singh JP, Fazel S. Forensic risk assessment: A metareview. *Criminal Justice and Behavior*. 2010; 37(9):965–988.
- Skeem JL. Risk technology in sentencing: Testing the promises and perils (Commentary on Hannah-Moffat, 2011). *Justice Quarterly*. 2013; 30(2):297–303.
- Skeem JL, Lowenkamp CT. Risk, race, and recidivism: Predictive bias and disparate impact. *SSRN Electronic Journal* 2016
- Skeem JL, Manchak S, Peterson JK. Correctional policy for offenders with mental illness: Creating a new paradigm for recidivism reduction. *Law and Human Behavior*. 2011; 35(2):110–126. [PubMed: 20390443]
- Skogan W, Frydlik K. *Fairness and Effectiveness in Policing: The Evidence* Washington, DC: Committee on Law and Justice; Division of Behavioral and Social Sciences and Education; National Research Council; National Academies Press; 2004
- Storey JE, Kropp PR, Hart SD, et al. Assessment and management of risk for intimate partner violence by police officers using the Brief Spousal Assault Form for the evaluation of risk. *Criminal Justice and Behavior*. 2014; 41(2):256–271.
- Summers C, Willis T. *Pretrial risk assessment: Research summary* Arlington, VA: Bureau of Justice Assistance, US Department of Justice; 2010
- Taxman FS, Marlowe D. Risk, needs, responsivity: In action or inaction? *Crime and Delinquency*. 2006; 52(1):3–6.
- Taxman FS, Thanner M, Weisburd D. Risk, Need, and Responsivity (RNR): It all depends. *Crime and Delinquency*. 2006; 52(1):28–51. [PubMed: 18542715]
- Trujillo MP, Ross S. Police response to domestic violence. *Journal of Interpersonal Violence*. 2008; 23(4):454–473. [PubMed: 18252942]
- Van Nostrand M, Keebler G. Pretrial risk assessment in the federal court. *Federal Probation*. 2009; 73(2) <http://ezproxy.cul.columbia.edu/login?url=http://search.proquest.com/ezproxy.cul.columbia.edu/docview/213983460?accountid=10226>.
- Weisburd D, Eck JE. What can police do to reduce crime, disorder, and fear? *The Annals of the American Academy of Political and Social Science*. 2004; 593(1):42–65.

WesternB. Punishment and Inequality in AmericaNew York: Russell Sage Foundation; 2006

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript




Figure 1. Directed acyclic graphs illustrating a scenario in which the predictive values of a relationship (i.e. between antisocial peers and re-arrest) may be identical, whereas optimal intervention points may differ (antisocial behavior in panel (a) and antisocial peers in panel (b)).




Figure 2. Directed acyclic graph illustrating a scenario in which a strong predictive association (i.e. between antisocial cognitions and re-arrest) has no causal value, because intensive policing is a complete confounder.